

Weighted Positive Binary Decision Diagrams for Exact Probabilistic Inference

Giso H. Dal and Peter J.F. Lucas,
Institute for Computing and Information Sciences,
Radboud University Nijmegen, The Netherlands
{gdal,peter1}@cs.ru.nl

Abstract

Recent work on weighted model counting has been very successfully applied to the problem of probabilistic inference in Bayesian networks. The probability distribution is encoded into a Boolean normal form and compiled to a target language, in order to represent local structure expressed among conditional probabilities more efficiently. We show that further improvements are possible, by exploiting the knowledge that is lost during the encoding phase and incorporating it into a compiler inspired by Satisfiability Modulo Theories. Constraints among variables are used as a background theory, which allows us to optimize the Shannon decomposition. We propose a new language, called *Weighted Positive Binary Decision Diagrams*, that reduces the cost of probabilistic inference by using this decomposition variant to induce an arithmetic circuit of reduced size.

Keywords: knowledge compilation, probabilistic inference, weighted model counting, Bayesian networks, binary decision diagrams.

1. Introduction

Bayesian networks, BNs for short, have been a subject of great interest partly due to their contribution in solving real-life problems that involve uncertainty. Bayesian networks are probabilistic graphical models that represent joint probability distributions concisely by factoring them into conditional probabilities based on independence assumptions, in order to perform inference more efficiently [1]. Further representational and computational advances have been made by exploiting causal independence [2], as well as contextual independence [3] and determinism [4] expressed in conditional probability tables (CPTs). In order to capture these independencies local to CPTs, Bayesian networks have been represented as weighted Boolean formulas [5, 6], reducing inference to *Weighted Model Counting* (WMC), or *weighted #SAT* [5]. By representing a Bayesian network as Boolean formula f in *conjunctive normal form* (CNF), it can be compiled into a more concise normal form, or language, that renders inference a polytime operation in the size of the representation [7].

A joint probability space with n Boolean variables has 2^n interpretations. It is therefore necessary to be able to reason with sets of interpretations, requiring a symbolic representation [8]. Symbolic inference unifies the work of probabilistic inference and the extensive research done in the field of model checking, verification and satisfiability [9]. *Ordered Binary Decision Diagrams* (OBDDs) are based on Shannon decompositions and have been a very

influential symbolic representation that reduces compilation to the problem of finding the variable ordering resulting in the optimal factoring.

We focus on the disadvantage of the approaches in recent work that encode a BN as an independent CNF f , motivated by the ability to use off-the-shelf SAT-solvers. While maintaining this ability, we exploit the knowledge that is lost during the encoding without requiring this independence.

Our contributions are the following. We propose a weighted variant of OBDDs, called *Weighted Positive Binary Decision Diagrams* (WPBDDs), which are based on *positive Shannon decompositions*, allowing constraints in BNs to be represented more concisely. We use probabilities as *symbolic* edge weights, reducing the search space exponentially. An optimized compilation algorithm is introduced, inspired by the field of Satisfiability Modulo Theories (SMT), namely a *lazy SMT-solver* [10]. It provides the means to view constraints among variables in the encoding as background theory \mathcal{T} which supports the SAT-solver, allowing *constant time conditioning*. We compile the conditional probability tables of a BN explicitly, but leave implicit the domain closure implied by the encoding. This approach allows us to remove by up to a third of the clauses in the encoding.

A comparison is provided with the state-of-the-art CUDD (CU Decision Diagram [11]) and SDD (Sentential Decision Diagram [12]) compilers and we show that WPBDDs induce arithmetic circuits that are 60% reduced in size *on average* compared to a corresponding OBDD circuits at representing over 30 publicly available BNs. We show an inference speedup of over 2.6 times on average compared to Weighted Model Counting with OBDDs, and an average speedup between 5 and 1000 times compared to different implementations of the Junction Tree algorithm, making WPBDDs a valuable addition to the field of exact probabilistic inference.

After preliminaries and background (Section 3), we introduce WPBDDs (Section 4). The process of using a BN to perform exact inference by WMC is explained (Section 5) in addition to its optimization (Section 6). We conclude with experimental results (Section 7), and review achievements (Section 8). First we summarize related work.

2. Related Work

Probabilistic inference is a hard computational problem that can be achieved by marginalizing out non-evidence variables from a joint distribution, requiring an exponential number of operations in the worst case. Efforts toward efficient exact probabilistic inference attempt to find a concise factorization, e.g., BNs represent a joint probability distribution as a multiplicative factorization, exploiting independencies among variables. Further improvements to this factorization have been made by using propositional and first order logic [13, 14, 15]. Representing probability distributions as Boolean functions empowers probabilistic inference with the tools developed for VLSI-CAD design, by using symbolic representations and Boolean algebra for minimization. *Symbolic Probabilistic Inference* (SPI) [9] is a good example of this, which is currently more commonly referred to as inference by WMC or #SAT [16].

A BN can be viewed as a constraint satisfaction problem (CSP) and translated to a satisfiability (SAT) instance in *conjunctive normal form* (CNF), a form commonly used in satisfiability solving. Various encodings of CSPs have been proposed, namely log, direct [17], order [18], compact order [19], log-support encoding [20], etc. In the context of BNs, a probability distribution can be considered a pseudo Boolean function $f : \{0, 1\}^n \rightarrow \mathbb{R}$, with arity n , which can be uniquely written as an exponentially sized multi-linear polynomial

[21, 22]. Others have used the direct encoding [6], or a combination between the direct and order encoding [23], where a BN is viewed as a set of discrete real valued functions, where each function represents a distinct CPT.

Inference by WMC is motivated by linear time complexity in the size of the representation [24], where the common goal is to exploit local structure [25]. Choosing a representation or *language* to compile to is therefore a critical task, where one must deal with the balance between the functions a language can represent concisely versus its algorithmic properties. Initial attempts include probability trees [3, 26, 27] and recursive (factored) probability trees [28], which focus on concisely representing each CPT independently, allowing their usage in inference algorithms directly. Probabilistic Decision Graphs (PDG) have even shown that the smallest PDG is at least as small as the smallest Junction tree for the same distribution [29].

Current WMC approaches to inference divide into *search* and *compilation* methods [30]. Typical search algorithms are based on DPLL-style SAT solvers that do an exhaustive run to count all satisfying models [23]. Recording SAT evaluation paths (i.e. resolution steps) as a compiled structure (e.g. an OBDD), yields one possible factoring. We refer to finding the optimal factoring given all variable orderings, as *exact* compilation.

Compilation performance has been improved by clause learning [31], formula caching [32], bounding [33], and using canonical languages. Representational advances include symmetry detection [34, 35], support for causal independence [36] and using *read-once functions* [37].

Representations relevant in the context of BN compilation are *AND/OR Multi-Valued Decision Diagrams* (AOMDD) [38], *Sentential Decision Diagrams* (SDD) [12], Zero-suppressed Binary decision diagrams (ZBDD) [21] and Ordered Binary Decision Diagrams (OBDD) [39], which view probabilities as auxiliary literals, resulting in an intractably large search space. Multi-Terminal BDDs [40] represent multi-valued functions, but would require too many terminal nodes considering the size of a probability distribution. Variants of Edge-Value BDDs [41, 42] focus on real valued functions. When multiple CPTs have probabilities in common, we lose the ability to distinguish from which CPT the probabilities originated. We therefore cannot determine on which variables they depend, resulting in an inconsistent model count with regard to the distribution. Our approach to maintain consistency is to represent probabilistic edges weights *symbolically*. This differentiates our approach from Multi-Terminal BDDs [40] and Edge-Value BDDs [41]. And unlike SDDs [43], we are not obligated to view probabilities as auxiliary literals, reducing the search space to a fraction of its former size. A common characteristic with ZBDDs, is the ability to represent mutual-exclusive constraints more concisely [44]. The intuitive difference is that ZBDD optimize only the positive cofactor, while we optimize both the positive and negative cofactor of decomposition nodes, a matter we will elaborate on in upcoming discussions.

3. Preliminaries and Background

We provide here a description of what Bayesian networks are and introduce a running example (Section 3.1), show how to encode a BN onto the Boolean domain (Section 3.2), and describe an influential representation that will serve for comparison with ours (Section 3.3).

3.1. Bayesian Networks

Probabilistic inference is an important computational problem in Artificial Intelligence. A full joint probability distribution defined over n Boolean variables is of size $\mathcal{O}(2^n)$. Finding

the minimal representation of a function describing a probability distribution reduces memory and inference complexity, which is the motivation for this paper.

A Bayesian Network (BN) is a graphical representation that is used to compactly represent a joint distribution as a product of factors, by taking advantage of *conditional independence* (CI). A BN is an directed acyclic graph (DAG) that models variables X as nodes, the dependencies among them as edges, and their joint probability distribution as

$$P(X) = P(x^1, \dots, x^n) = \prod_{i=1}^n P(x^i \mid pa(x^i)),$$

where $P(x^i \mid pa(x^i))$ represents the conditional probability of variable x^i given its parents $pa(x^i)$. Conditional probability tables (CPTs) are associated with edges and capture the degree to which variables are related. BNs reduce the size of representing a probability distribution to $\mathcal{O}(n2^k)$, where k is the maximum number of parents of any node.

Example 1. Figure 1 shows BN \mathcal{B} defined over variables $X = \{a, b\}$ (Figure 1b), its CPTs (Figure 1c) and corresponding full joint probability distribution (Figure 1a).

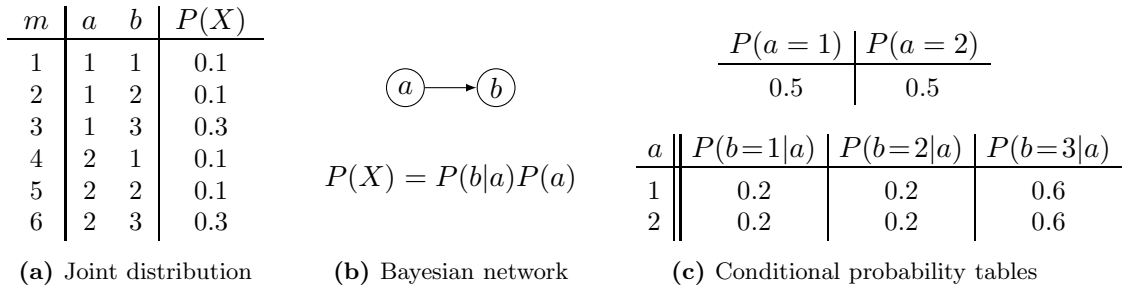


Figure 1 Bayesian network with local structure

Figure 1b includes a factored form that can greatly be improved when CPTs exhibit *local structure*, which comes in two forms. *Context-specific independence* (CSI) is expressed when probabilities in a CPT show uniformity regardless of the value of one or more variables they have in common, with or without a certain context. *Determinism* is expressed when probabilities in a CPT are equal to 0 or 1, which can be used to simplify the Boolean formula that represents it.

In order to exploit more of the problem structure than a BN, we *compile* it to a target language that is more capable of doing so. The compilation process is not just about reformulating into a different language, it is also about finding the minimal representation given that language. The goal is to derive an arithmetic circuit corresponding to that representation with improved complexity compared to the standard factorization.

3.2. Encoding Bayesian Networks

In order to exploit local structure, we encode BNs into *conjunctive normal form* (CNF), which is the most common representation used in satisfiability solving. It consists of a conjunction of clauses, where each clause is a disjunction of literals. A literal is a propositional Boolean variable or its negation. A Bayesian Network defined over variables X can be seen as a multi-linear function $f : X \rightarrow \mathbb{R}$. Using a weighted adaptation of the *sparse* or *direct*

encoding [6, 17, 45], we encode f into a Boolean function $\mathcal{E}(f) = f^e$ by representing it as a weighted CNF:

$$\mathcal{E}(f) = f^c \wedge f^m, \quad (1)$$

where constraint clauses f^c support the mapping $\mathcal{M}(f) = f^m$ that encodes probabilities and introduces a Boolean variable for each unique variable-value pair. Details are discussed below.

3.2.1. Encoding Constraints

The mapping function \mathcal{M} introduces for each $x \in X$ atoms $\mathcal{A}(x) = \{x_1, \dots, x_n\}$, where x_i signifies x being equal to its i^{th} value. To maintain consistency among variables we add to f^c an *at-least-once* (ALO) constraint clause for each x , to ensure x is assigned a value:

$$(x_1 \vee \dots \vee x_n) \quad (2)$$

As values of a variable are mutually exclusive, we add to f^c the following *at-most-once* (AMO) constraint clauses:

$$\bigwedge_{i=1}^n \left(x_i \implies \bigwedge_{x_j \in \mathcal{A}(x) \setminus \{x_i\}} \overline{x_j} \right) = \bigwedge_{i=1}^n \bigwedge_{j=i+1}^n (\overline{x_i} \vee \overline{x_j}), \quad (3)$$

where $\overline{x_i}$ indicates the negation of x_i .

3.2.2. Encoding CPTs

The BN's factored form is preserved by using a weighted adaptation of the *direct* encoding. The mapping function \mathcal{M} adds a clause for every probability $P(x|U)$, where x depends on variables $U = \{u^1, \dots, u^r\}$:

$$(x \wedge u^1 \wedge \dots \wedge u^r \implies \omega_i) = (\overline{x} \vee \overline{u^1} \vee \dots \vee \overline{u^r} \vee \omega_i), \quad (4)$$

which we henceforth shall view as a *weighted clause* $(\overline{x} \vee \overline{u^1} \vee \dots \vee \overline{u^r}) : \omega_i$, where its symbolic weight ω_i represents probability $P(x|U)$. We introduce a symbolic weight for every unique probability per CPT, and thus allow multiple models to be associated with it:

$$I^1 \wedge \dots \wedge I^s, \quad (5)$$

where each I^j has weight ω_i and forms the *implicate* of model m associated with $P(\mathbf{x}|U)$, i.e., $I^j = \omega_i : (\mathbf{x} \wedge \mathbf{u}^1 \wedge \dots \wedge \mathbf{u}^r)$, where the variables $\mathbf{x} \cup U$ have been appropriately mapped by \mathcal{M} , i.e., $\mathbf{x} \in \mathcal{A}(x)$ and each $\mathbf{u}^k \in \mathcal{A}(u^k)$.

Example 2. Assume a BN as given in Example 1. The following clauses form f^c :

Variable	ALO (Eq. 2)	AMO (Eq. 3)
a	$(a_1 \vee a_2)$	$(\overline{a_1} \vee \overline{a_2})$
b	$(b_1 \vee b_2 \vee b_3)$	$(\overline{b_1} \vee \overline{b_2}) \wedge (\overline{b_1} \vee \overline{b_3}) \wedge (\overline{b_2} \vee \overline{b_3})$

The variables that make up the search space during compilation therefore are $\mathcal{A}(\{a, b\}) = \{a_1, a_2, b_1, b_2, b_3\}$. We encode equal probabilities $P(x|U)$ as unique symbolic weights per CPT:

$P(a_1)$	$P(a_2)$	a	$P(b_1 a)$	$P(b_2 a)$	$P(b_3 a)$
ω_1	ω_1	1	ω_2	ω_2	ω_3
		2	ω_2	ω_2	ω_3

In accordance with Equations 4 and 5, f^m consists of the following clauses, accompanied by their respective symbolic weights:

$$\begin{aligned}
& (\overline{a_1}) : \omega_1 \wedge (\overline{a_2}) : \omega_2 \wedge \\
& (\overline{a_1} \vee \overline{b_1}) : \omega_2 \wedge (\overline{a_1} \vee \overline{b_2}) : \omega_2 \wedge (\overline{a_1} \vee \overline{b_1}) : \omega_2 \wedge \\
& (\overline{a_1} \vee \overline{b_2}) : \omega_2 \wedge (\overline{a_1} \vee \overline{b_3}) : \omega_3 \wedge (\overline{a_2} \vee \overline{b_3}) : \omega_3.
\end{aligned}$$

3.3. Ordered Binary Decision Diagrams

A Boolean function f defined over a set of variables X is a function that maps each complete assignment of its variables to either *true* (1) or *false* (0). The *conditioning* of f on instantiated variable x_i is defined as the projection:

$$f_{|x_i \leftarrow b}(x_1, \dots, x_n) = f(x_1, \dots, x_{i-1}, b, x_{i+1}, \dots, x_n), \quad (6)$$

with $b \in \{0, 1\}$. We will use shorthand notations $f_{|x_i}$ and $f_{|\overline{x_i}}$ for $f_{|x_i \leftarrow 1}$ and $f_{|x_i \leftarrow 0}$, respectively. Shannon's theorem is used to find a more compact way to represent f by *factoring* it.

Theorem 1. [46] Shannon's expansion allows a Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ defined over variables X , to be written in terms of its inputs:

$$f = x \wedge f_{|x} \vee \overline{x} \wedge f_{|\overline{x}},$$

with $x \in X$ and where $f_{|x}$ is called the *positive cofactor* of f with respect to x , and $f_{|\overline{x}}$ the *negative cofactor*. Applying Shannon's theorem is known as an (additive) *decomposition step*. The decomposition of f is defined as the recursive application of Shannon's expansion theorem to cofactors, removing one variable at a time, until no variables are left.

Note that all proofs of theorems, lemmas, etc., can be found in the Appendix. The *Shannon decomposition* is key to one of the most influential representations in Artificial Intelligence (AI), namely *Ordered Binary Decision Diagrams* (OBDD)[47].

Definition 1. [48] A Binary Decision Diagram (*BDD*) represents Boolean function f defined over variables X as a rooted, directed acyclic graph, where each node v represents a Shannon decomposition on variable $\text{var}(v) \in X$. A BDD is *ordered* (*OBDD*) if variables appear in the same order on all paths from the root. It is a *canonical representation* if it is reduced by applying the following rules:

1. *Merge rule:* All isomorphic subgraphs are merged.
2. *Delete rule:* All nodes are removed whose children are isomorphic.

One can compile the described encoding of BNs to OBDDs to perform inference by WMC. Although other languages have been used in this context, we focus more on OBDDs as it is commonly used in comparisons with related work.

4. Weighted Positive Binary Decision Diagrams

Consider variable x and its corresponding mapped atoms $\mathcal{A}(x) = \{x_1, x_2\}$. Decision diagrams that are based on Shannon decompositions produce an unnecessarily large arithmetic circuit by redundantly representing constraints f^c provided as part of encoding \mathcal{E} . They also do not take advantage of the symmetric relation $x_1 = \overline{x_2}$ and $\overline{x_1} = x_2$ in the presence of ALO constraints. To ameliorate this, we propose a new canonical language called *Weighted Positive Binary Decision Diagrams* (WPBDD), which are based on *positive Shannon decompositions* and *implicit conditioning*. We will elaborate on these concepts through an intermediate unweighted variant of WPBDDs (PBDD).

4.1. Explicit and Implicit Conditioning

When encoded function f^e contains a *unit clause* x_i (a clause consisting of a single literal), it can be simplified using *unit propagation*:

1. Every clause containing x_i is removed, excluding the unit clause.
2. Literal $\overline{x_i}$ is removed from every clause containing it.

When conditioning f^e on literal x_i we are guaranteed, due to constraints f^c , to obtain unit clauses containing negated literals $\overline{x_j}$, with $x_j \in \mathcal{A}(x) \setminus x_i$ (i.e., if x is equal to its i^{th} value, it cannot be equal to its j^{th} value).

Example 3. We will show by example what unit clauses are obtained by conditioning on positive literals. Consider the constraint clauses provided by Example 2, regarding only variable b :

$$f^c = (b_1 \vee b_2 \vee b_3) \wedge (\overline{b_1} \vee \overline{b_2}) \wedge (\overline{b_1} \vee \overline{b_3}) \wedge (\overline{b_2} \vee \overline{b_3})$$

According to Shannon's expansion the following holds:

$$f^c = b_1 \wedge f_{|b_1}^c \quad \vee \quad \overline{b_1} \wedge f_{|\overline{b_1}}^c$$

Now specifically look at the unit clauses that result from conditioning on positive literal b_1 , i.e., instantiating b_1 and performing unit propagation.

$$\begin{aligned} f_{|b_1}^c &= (1 \vee b_2 \vee b_3) \wedge (0 \vee \overline{b_2}) \wedge (0 \vee \overline{b_3}) \wedge (\overline{b_2} \vee \overline{b_3}) \\ &= (1) \wedge (\overline{b_2}) \wedge (\overline{b_3}) \wedge (\overline{b_2} \vee \overline{b_3}) \\ &= (1) \wedge (\overline{b_2}) \wedge (\overline{b_3}) \\ &= (\overline{b_2}) \wedge (\overline{b_3}). \end{aligned}$$

Thus, conditioning on b_i will result in unit clauses containing negated literals $\overline{b_j}$, with $b_j \in \mathcal{A}(b) \setminus \{b_i\}$.

We distinguish between two types of conditioning based on the previous observation, and describe them in the following definition.

Definition 2. Let f^e be an encoded representation of function f , given encoding \mathcal{E} , where f is defined over variables X . We define $f_{\parallel x_i}^e$ as the conditioning of f^e on literals $\{x_i, \overline{x_1}, \dots, \overline{x_{i-1}}, \overline{x_{i+1}}, \dots, \overline{x_n}\}$ in any order, i.e., as the explicit conditioning of f^e on literal $x_i \in \mathcal{A}(x)$, and its implicit conditioning on literals $\overline{x_j} \in \mathcal{A}(x) \setminus x_i$, with $x \in X$:

$$f_{\parallel x_i}^e = f_{|x_i, \overline{x_1}, \dots, \overline{x_{i-1}}, \overline{x_{i+1}}, \dots, \overline{x_n}}^e,$$

It follows from Definition 2 and the constraints provided by encoding \mathcal{E} , that the relation between $f_{\parallel x_i}^e$ and $f_{|x_i}^e$ is given by the following equality:

$$f_{|x_i}^e = \left(\bigwedge_{\overline{x_j} \in \mathcal{A}(x) \setminus x_i} \overline{x_j} \right) \wedge f_{\parallel x_i}^e. \quad (7)$$

Implicit conditioning on unit clauses takes advantage of deterministic behavior expressed in f^c , while other representations would have to explicitly condition these unit clauses out. The advantage is two-fold. The size of the encoding can be reduced by removing constraint clauses f^c generated by Equation 2 and 3, and integrating them directly into the compilation process through theory \mathcal{T} . As will be shown later, by separating constraint clauses from f^e as theory \mathcal{T} allows them to be conditioned in *constant time*, as opposed to quadratic time. Secondly, the size of the compiled structure is reduced by not having to represent redundant constraint information with our variant on the Shannon expansion, introduced in the following section.

4.2. Positive Shannon Decomposition

We propose *positive Shannon decompositions* that use background knowledge to improve upon Shannon decompositions by combining it with implicit conditioning.

Lemma 1. A positive Shannon expansion allows an encoded Boolean function $\mathcal{E}(f) = f^e$, where f is defined over X , to be written in terms of its inputs:

$$f^e = f^c \wedge \left(x_i \wedge f_{\parallel x_i}^e \vee f_{\overline{x_i}}^e \right),$$

where $f^e = f^c \wedge f^m$, and $x_i \in \mathcal{A}(x)$, with $x \in X$. The positive cofactor $f_{\parallel x_i}^e$ incorporates implicit conditioning (Definition 2). The negative cofactor $f_{\overline{x_i}}^e$ and the decomposition of f^e are as defined by Theorem 1.

As intuition might confirm, logical representations will grow by the reintroduction of constraints at every expansion. We introduce a *reduced* expansion by removing constraint clauses f^c that introduces additional models, in turn allowing us to find more concise representations. These models can easily be removed by a post decomposition conjoin with f^c .

Theorem 2. A reduced positive Shannon expansion allows an encoded Boolean function $\mathcal{E}(f) = f^e$, where f is defined over X , to be written in terms of its inputs under constraints f^c :

$$f^e \models x_i \wedge f_{\parallel x_i}^e \vee f_{\overline{x_i}}^e,$$

where $f^e = f^c \wedge f^m$, and $x_i \in \mathcal{A}(x)$, with $x \in X$. Cofactors and decomposition are as defined by Lemma 1.

By using the reduced form of the expansion, we are able represent constraints more concisely in corresponding logical circuits (Figure 2), as well as in the soon to be introduced representation that utilizes it.




Figure 2 Logical Circuits

Using encoding \mathcal{E} we can infer that the delete rule used to reduce OBDDs will never be applied to constraint clauses f^c . OBDDs are therefore not capable of capturing local structure along one dimension. To ameliorate this, we introduce *positive* OBDDs (PBDD) as an unweighted intermediate representation, that are based on the positive Shannon decomposition and substitutes the delete rule with the *collapse* rule, which as opposed to deleting literals, applies the *distributive law* to involved literals in the induced logical circuit.

Definition 3. A positive OBDD (PBDD) represents Boolean function $\mathcal{E}(f) = f^e$, where f is defined over variables X , as an ordered BDD where each node v represents a positive Shannon decomposition on variable $\text{var}(v) \in \mathcal{A}(X)$. It is a canonical representation if reduced by applying the following rules:

1. *Merge rule:* All isomorphic subgraphs are merged.
2. *Collapse rule:* remove direct descendant u of node v iff $f|_{x_i} = f|_{x_j}$, where $\text{var}(v) = x_i$ and $\text{var}(u) = x_j$, with $x_i, x_j \in \mathcal{A}(x)$ and $x \in X$.

A function *essentially* depends on a variable if it appears in its prime implicate. The variable set \mathcal{S} , on which f^e essentially depends, is called the *support* of f^e . We will use this support set to identify to what variables the collapsed rule has been applied in order to produce its corresponding arithmetic circuit, a trick similarly utilized with Zero-Suppressed BDDs (ZBDD) [44]. Note that a Boolean function $\mathcal{E}(f) = f^e$, where f is defined over variables X , essentially depends on $\mathcal{A}(X)$, because f^c is a prime implicate that mentions all $\mathcal{A}(X)$. The canonical property of PBDDs follows from the fact that a binary tree can be reconstructed from a PBDD and its support set, by apply its reduction rules reversely.

Figure 3 shows the difference in representational size between an OBDD, a ZBDD and a PBDD representing the same function with constraints on $\mathcal{A}(x) = \{x_1, x_2\}$, where g is a Boolean function that does not essentially depend on literals $\{x_1, x_2\}$. The positive Shannon decomposition does not only reduce the size the corresponding logical circuit, it also reduces the size of the representation. More generally, OBDDs require an exponential number of nodes in the product of each constraint variable's dimension, where ZBDDs require a linear number, and where PBDDs require only 1 node.

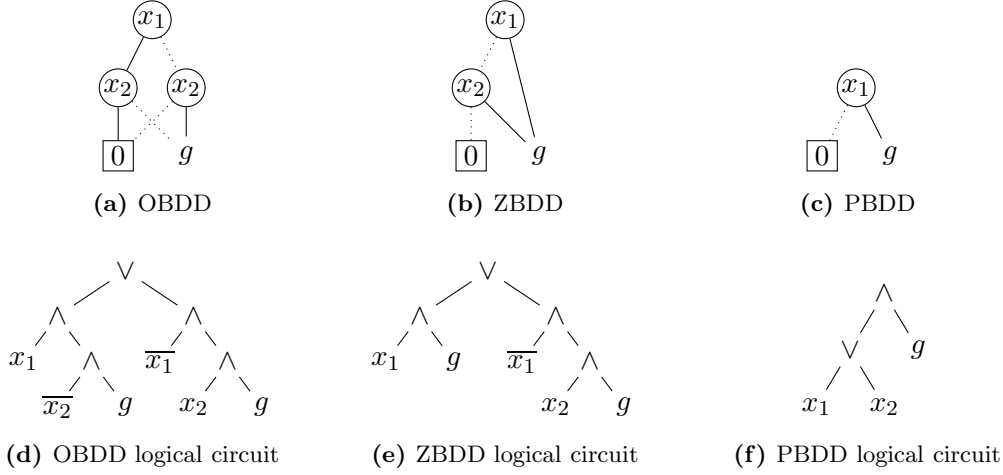


Figure 3 From BDDs to logical circuits

Figure 3 shows that there are functions where the corresponding minimal PBDD and OBDD differ exponentially in size. Not represented in the figure, the collapse rule additionally removes the nodes that share the same positive cofactor with their parents (Figure 4).




Figure 4 Application of the collapse rule, where cofactors are equal

The semantics of nodes whose child has been removed changes. A missing node, inferable by the support set \mathcal{S} and variable order, indicates the application of the collapse rule, i.e., the distributive law on x_i and x_j . There is no ambiguity regarding the delete rule as it can never be applied on $\mathcal{A}(X)$ due to constraint clauses f^c . Note that the delete rule can however be applied in case one optimizes Boolean variables of the BN by representing them with only one literal in f^c , as opposed to two. This will delete literals from the induced circuit and result in an inconsistent model count with regard to the probability distribution. It is precisely for this reason why the collapse rule uses the distributive law on involved literals to simplify the induced circuit, as opposed to deleting them from it. The combination of the merge and collapse rule allow for more fine grained control in exploiting CSI, because it allows independence given a subset of the values to be expressed more efficiently when dealing with multi-valued variables.

Proposition 1. *An OBDD representing Boolean function f and an PBDD representing $\mathcal{E}(f)$ induce isomorphic logical circuits under Boolean identity, given an appropriate ordering.*

Proposition 2. *Given an ordering on $\mathcal{A}(X)$, the size of PBDD φ is less than the size of OBDD ψ when both representing $\mathcal{E}(f) = f^c$, where f is defined over variables X .*

4.3. Adding Probabilities as Weights

Encoding \mathcal{E} represents a BN as a weighted propositional formula. We extend PBDDs to *weighted* PBDDs (WPBDD) using an intuitive scheme, taking advantage of the fact that probabilities are fully implied by the variables in the BN. Traditionally, an empty clause would result in a contradiction, i.e., the instantiation is unsatisfiable. We *implicitly* assign weight ω of empty clause c to the edge it is associated with, and remove c from the expression. When multiple empty clauses are associated with an edge, we simply assign the conjunction (multiplication) of their weights to the edge. To maintain canonicity, we only assign weights on the side of the positive cofactor.

The collapse rule previously introduced can easily be extended for the non-binary and weighted case as shown in Figure 5.




Figure 5 Application of the collapse rule, where functions $f_{||x_i} = f_{||x_j}$ and weights $\omega_i = \omega_j$

Definition 4. A weighted PBDD (WPBDD) representing Boolean function $\mathcal{E}(f) = f^e$, where f is defined over variables X , is a PBDD where each node v is a tuple $\langle x_i, W, f_{||x_i}^e, f_{|\overline{x_i}}^e \rangle$ that represents a weighted reduced positive Shannon expansion:

$$f^e \models x_i \wedge (W \wedge f_{||x_i}^e) \vee f_{|\overline{x_i}}^e,$$

where $x_i \in \mathcal{A}(X)$, and W is a conjunction of weights ω_i that correspond to $f_{||x_i}^e$. Positive and negative cofactors are as described by Lemma 1. It is a canonical representation if reduced by applying the following rules:

1. Merge rule: All isomorphic subgraphs are merged.
2. Collapse rule: remove direct descendant u of node v iff $W \wedge f_{||x_i} = W \wedge f_{||x_j}$, where $\text{var}(v) = x_i$ and $\text{var}(u) = x_j$, with $x_i, x_j \in \mathcal{A}(x)$ and $x \in X$.

Example 4. Consider the CPTs from Example 1, where per CPT, equal probabilities are represented by unique symbolic weights ω_i .

$P(a_1)$	$P(a_2)$	a	$P(b_1 a)$	$P(b_2 a)$	$P(b_3 a)$
$\{\omega_1\}$	$\{\omega_1\}$	1	$\{\omega_2\}$	$\{\omega_2\}$	$\{\omega_3\}$
		2	$\{\omega_2\}$	$\{\omega_2\}$	$\{\omega_3\}$

Figure 6 shows the minimization of a WPBDD using variable ordering $a_1 < a_2 < b_1 < b_2 < b_3$, that represents the BN with 3 probabilities instead of 8. Figure 7 shows the comparison of this WPBDD with an OBDD representing the same function, given variable ordering $a_1 < a_2 < \omega_1 < b_1 < b_2 < b_3 < \omega_2 < \omega_3$, which results in a minimal OBDD that obeys the partial ordering used for the WPBDD.




Figure 6 WPBDD reduction

5. Symbolic Inference

We perform Bayesian inference through a three phase process: encoding, compiling and model counting:

	Composition	Input	Output
Encoder	Encoding \mathcal{E}	f	Theory $\mathcal{T} + f^e$
Compiler	\mathcal{T} -solver + SAT	$\mathcal{T} + f^e$	WPBDD φ
Counter	\mathcal{T} -solver + WMC	$\mathcal{T} + \varphi$	$P(x e)$

A BN represented by f is first encoded by the *encoder* as Boolean function f^e using encoding \mathcal{E} as defined in Section 3.2. The encoder also provides background theory \mathcal{T} representing f^c , i.e., the constraints among variables that support mapping \mathcal{M} .

The *compiler* uses a lazy SMT-solver to record evaluation paths and as a WPBDD, given f^e and theory \mathcal{T} . A lazy SMT-solver combines a SAT-solver with a theory-solver (or \mathcal{T} -solver) for some theory \mathcal{T} . Traditionally, the role of a theory-solver is to purely report back on the satisfiability of \mathcal{T} . We have extended the theory-solver to provide more information in order to support implicit conditioning in the SAT-solver.




Figure 7 WPBDD and OBDD comparison

The *counter* computes the probability of x given evidence \mathbf{e} , by translating the provided WPBDD into an arithmetic circuit and using the extended capabilities of the T -solver to properly instantiate the variables.

5.1. Compilation

The order of decomposition greatly influences representation size. Compilation therefore reduces to finding the optimal variable ordering. Satisfiability (SAT) is key during compilation. Normally, when CNF f contains an empty clause we derive a contradiction. Note that if all contradicting clauses are weighted, we supersede the contradiction and introduce their weights into our representation at the corresponding edge. In order to obtain a minimal BDD representation, the search space of all variable instantiations is traversed with a DPLL-style algorithm, in order to find partial instantiations that describe equal (sub)functions [5], as depicted in Figure 8.

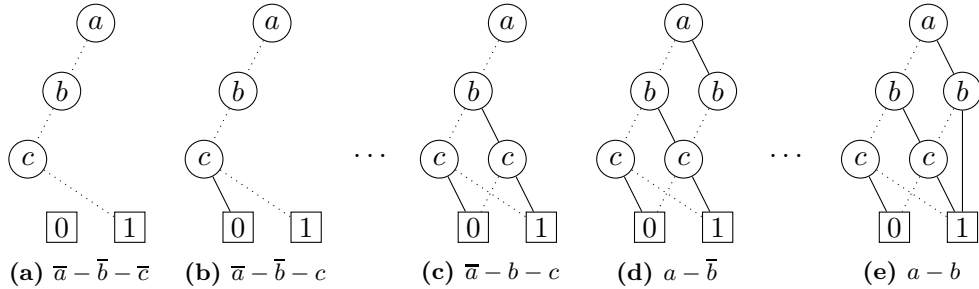


Figure 8 Compilation example for non-trivial function, given ordering $a \rightarrow b \rightarrow c$

Algorithm 1 shows a Depth-first search/dynamic programming (DFS+DP) approach that uses a lazy SMT-solver (*solver*) to compile f^e into a WPBDD, given some ordering on the variables (*ordering*), where \top and \perp denote the terminal nodes representing *true* and *false*, respectively.

The lazy SMT-solver is unique in the sense that there exists a link beyond satisfiability feedback between the \mathcal{T} -solver and the SAT-solver (*solver.theory* and *solver.sat*, re-

Algorithm 1 Compiler

```

1  struct node {
2      literal l;
3      node *t, *e;
4      set W;
5  };
6
7  enum satisfy_t {
8      satisfiable = 0,
9      unsatisfied = 1,
10     unsatisfiable = 3
11 };
12
13 satisfy_t solver::condition(literal l){
14     solver.theory.condition(l);
15     solver.sat.condition(l);
16     solver.sat.condition(
17         solver.theory.unit_clauses());
18
19     return solver.theory.state() |
20         solver.sat.state();
21 }
22
23 node* condition(literal l)
24     solver.condition(l);
25     if(not negated(l))
26         n->W += solver.sat.weights();
27
28     node *n;
29     switch(solver.state()){
30         case unsatisfied:
31             n = compile(new node, i+1);
32             break;
33         case satisfiable:
34             n =  $\top$ ;
35             break;
36         case unsatisfiable:
37             n =  $\perp$ ;
38             break;
39     }
40     solver.undo();
41
42     return n;
43 }
44
45 node* compile(n, i=0){
46     n->l = ordering[i];
47     n->t = condition(n->l);
48     n->e = condition(not n->l);
49
50     apply_collapse_rule(n);
51     apply_merge_rule(n);
52
53     return n;
54 }
55
56 wpbdd* compiler( $\mathcal{T}, f^e$ ){
57     solver.theory.init( $\mathcal{T}$ );
58     solver.sat.init( $f^m$ );
59
60     return compile(new node);
61 }

```

spectively). The traditional role of the theory-solver is to solely report back on satisfiability. To implement implicit conditioning, we have extended it to also provide its unit clauses, that are used by the SAT-solver to further condition f^e . This connection is possible because f^e and theory \mathcal{T} both essentially depend on the same variables. The SMT-solver reports f^e to be satisfiable only when both the theory and SAT-solver agree on this. Storing intermediate states as an undo mechanism for the solver is infeasible, thus it has the ability to dynamically undo any conditioning (`solver.undo`).

The compiler uses the satisfiability state of the SMT-solver (`solver.state`) to build the WPBDD and achieves a canonical form by applying to each subfunction the merge rule (as describe by [49]) and collapse rule (`apply_merge_rule(n)` and `apply_collapse_rule(n)`, respectively).

5.2. Inference by Weighted Model Counting

In order to perform inference by WMC, a WPBDD must be converted into a logical (refactored) form using Definition 4. Recall that a missing variable along a path implies the use of the distributive law, identifiable by using the variable ordering and support set \mathcal{S} . The logical form can easily be translated into an arithmetic circuit according to Table 1. Note that $x \vee y$ reduces to $x + y$, when x and y originate from the same dimension, i.e., $x, y \in \mathcal{A}(z)$, with $z \in X$.

Logical	Arithmetic
x	x
\overline{x}	$(1 - x)$
$x \wedge y$	$x * y$
$x \vee y$	$x + y - x * y$

Table 1 Converting logical to arithmetic operator

One of the reasons for using the positive Shannon decomposition is to prevent constraints among variables to be represented twice in the described process of symbolic inference: once as part of the compiled representation, and again when we substitute literals with their appropriate weight in order to perform model counting. During this later phase, theory \mathcal{T} is used to prohibit inconsistent network instantiations, preventing a state where multiple values are assigned to one variable. To perform inference, all weights ω_i are set to the probability they represent, and all other literals are set to 1. By conditioning \mathcal{T} on the evidence using the theory-solver, literals are found that conflict with the evidence in the form of unit clauses. These must be set to 0.

Example 5. Let $f^c = (a_1 \vee a_2) \wedge (\overline{a_1} \vee \overline{a_2})$ represent the constraint clauses for variable a of Example 1. When computing $P(a_1)$ we condition f^c on evidence a_1 yielding $f_{a_1}^c = \overline{a_2}$, thus evidence a_1 implies $a_1 = 1$ (true) and $a_2 = 0$ (false). This process is shown in Figure 9.




Figure 9 Probabilistic inference

6. Optimizations

6.1. Encoding

The constraint clauses f^c of encoded function $\mathcal{E}(f) = f^e$, where f is defined over variables X , introduce predictable symmetries into the encoding (demonstrated by Example 3). By incorporating these constraints directly into the compilation process through theory \mathcal{T} , the constraint clauses f^c generated by Equation 2 and 3 become obsolete and can thus be removed from f^e . This reduces the number of clauses in the encoding by:

$$\sum_{x \in X} \underbrace{1}_{\text{ALO clause}} + \underbrace{\binom{n}{2}}_{\text{AMO clauses}},$$

where we sum over every $x \in X$, with n the domain size of x , i.e., $|\mathcal{A}(x)|$. Both the at-least-once (ALO) and at-most-once (AMO) clauses contribute to reducing the number of clauses in the encoding to the number of probabilities in the CPTs of the BN. This gives an advantage over related work using the direct encoding, as it puts less strain on the SAT-solver by requiring it to only process $\mathcal{M}(f)$.

6.2. Compiler

The compiler uses a lazy SMT-solver, consisting of a theory- and SAT-solver. In the way we have build the compiler, it naturally allows for optimization by providing the ability to substitute the SAT-solver with any other state-of-the-art solver. We have optimized the SMT-solver by using the structure expressed by the encoding, and the fact that theory \mathcal{T} and f^e are defined over the same variables.

We have optimized the theory-solver such that it now supports *constant time conditioning* of constraint clauses, which can take up one third of the encoding as shown by experimental results later. All one needs is the function $V : \mathcal{A}(X) \rightarrow X$ that maps literal x_i back to x , where $x_i \in \mathcal{A}(x)$. For each x we maintain a counter that is initialized to the domain size of x , i.e., $|\mathcal{A}(x)|$. Conditioning on negated literal \bar{x}_i will decrease the counter corresponding to x by 1. If the counter reaches 0, we derive contradiction (i.e., unsatisfiable as x has no value). Conditioning on positive literal x_i will cause any following conditioning on $x_j \in \mathcal{A}(x) \setminus x_i$ as redundant (i.e., x can only have one value). The SAT-solver will be bypassed completely as a result and the compiler will continue with the next variable in the ordering, saving additional time.

We have simplified the SAT-solver considerably by taking advantage of the structure of $\mathcal{M}(f)$. We use an one-to-many map $Q : \mathcal{A}(X) \rightarrow \mathcal{O}$ from literal $l \in \mathcal{A}(X)$ to the clauses \mathcal{O} it occurs in, i.e., $Q(l) = \{c^1, \dots, c^n\}$. For each clause c^i , we maintain if it is satisfied with a counter, initialized to the number of literals it consists of. When conditioning on positive literal l we decrease counters associated with $Q(l)$ by 1. The clauses of which the counters have reached 0 are marked as satisfied, and their corresponding weights are set aside to be introduced into the representation later. Conditioning on negated literal \bar{l} will mark clauses $Q(\bar{l})$ as satisfied. This is possible because $\mathcal{M}(f)$ only contains negated literals, and we are able to assume that $Q(l) \cap Q(\bar{l}) = \emptyset$. When all clauses in f^m are satisfied, we derive f to be satisfiable given the evaluated instantiation. In combination with the SAT-solver being bypassed in the case of the previously mentioned redundant variables, this allows for conditioning in *linear time*, in the number of clauses that l occurs in.

7. Experimental Results

We have developed a tool chain, that can encode a Bayesian network into CNF, compile it to various different representations, and perform inference using the arithmetic circuits they induce. Using over 30 publicly available Bayesian networks, we provide empirical results on encoding size, representation size and compilation time comparisons to other well known

Bayesian Network	X	$\mathcal{A}(X)$	C	P	P^u	P^d
example	2	5	6	8	3	5
cancer	5	10	10	20	20	0
earthquake	5	10	10	20	20	0
asia	8	16	16	36	27	9
survey	6	14	16	37	37	0
student farm	12	25	26	70	44	26
sachs	8	24	32	228	175	53
poker	7	43	145	748	71	677
child	20	60	93	344	161	183
carpo	54	122	139	554	246	308
powerplant	40	120	160	432	360	72
alarm	37	105	143	752	182	570
win95pts	76	152	152	1148	274	874
insurance	27	89	142	1419	427	992
andes	220	440	440	2308	652	1656
hepar2	70	162	190	2139	1922	217
hailfinder	56	223	470	3741	835	2906
pigs	441	1323	1764	8427	1474	6953
link	714	1793	2304	20462	1282	19180
water	32	116	188	13484	3578	9906
munin1	186	992	3522	19226	4323	14903
pathfinder	135	520	3600	106432	2379	104053
weeduk	15	90	347	22611	4600	18011
fungiuk	15	165	1144	43007	8990	34017
munin2	1003	5376	19460	83920	23228	60692
munin3	1041	5601	20292	85615	24495	61120
munin	1041	5651	20432	98423	24222	74201
munin4	1038	5645	20426	97943	24621	73322
mildew	35	616	17550	547158	14772	532386
mainuk	48	421	3607	130180	18883	111297
diabetes	413	4682	31738	461069	17888	443181
barley	48	421	3607	130180	36924	93256

Table 2 Various statistics on BNs and their encoding, where the number of variables X , literals $\mathcal{A}(X)$, constraint clauses C , probabilities P , cumulative amount of unique probabilities per CPT P^u and the number of deterministic probabilities P^d are shown.

representations and compilers. We also compare the time it takes to perform exact inference compared to the classic Junction tree algorithm.

Statistics related to the encoding are shown in Table 2, which include Example 1 as BN `example`. The number of clauses produced by encoding \mathcal{E} is equal to $|f^c| + |P|$, for constraint clauses f^c and mapping \mathcal{M} , disregarding determinism. We can reduce the size of the encoding by up to a third, by moving constraint clauses to the theory solver, additionally allowing us to perform constant time conditioning on them. We can also see that the majority of the BNs will benefit greatly by the techniques in this paper by looking at the amount of equal and deterministic probabilities they contain.

We have developed a compiler that supports compilation of Bayesian networks to OBDDs

and ZBDDs (using the CUDD¹ 3.0.0 library), SDDs [43] (using the SDD² 1.1.1 library), and WPBDDs. Each decision diagram is created with the same ordering, within the same framework, i.e., doing the same amount of work in the same order. Quite literally, the only differences are the inserted appropriate function calls to different libraries, and the output representation. This will have comparative implications to whether a particular compilation will succeed given resource constraints as time and memory. At the same time, we did not tune the algorithms to ensure that our algorithm stood out favorably, ensuring fair comparison.

The framework divides the compilation process in two for efficiency. The logical representation of each CPT is first compiled separately, and then conjoined to represent the full distribution. The later is essential for producing a logical circuit with a consistent model count in order to perform for inference. All results regarding the WPBDD compiler have been produced with a hybrid approach, where CPTs are compiled in a topdown fashion, and conjoined bottom-up. We found that a fully bottomup approach is only favorable when BNs have large CPTs like `mildew`, where we got a 5x speedup compared to the topdown approach. In practice, large CPTs are usually avoid as they increase the complexity of inference.

Many strategies were explored in order to find a good variable ordering for each BN. Using simulated annealing in combination with an upper bound function yielded best results by far. The variable orderings were used to induce orderings based on literals, by saying that literal x_1 must come before y_1 if variable x comes before y in the variable ordering, where $x, y \in X$, $x_1 \in \mathcal{A}(x)$ and $y_1 \in \mathcal{A}(y)$. The weights are introduced into the ordering as literals precisely when the WPBDD would introduce them as edge weights.

Tables 3 and 4 show a comparative study between representation size and compilation time of supported representations, where WPBDD^{nc} is a WPBDD where the collapse rule has not been applied, in order to show the impact that this rule has. SDDs and SDD^r s are compiled using a balanced and right-aligned vtree ordering, respectively. A left-to-right traversal of these vtrees produces the ordering also used for the other representations. Table 3 indicates compilation failure due to a 24Gb RAM memory limitation or an one hour time limit by symbols - and *, respectively. The progress each failed compilation made before is indicated in Table 4. All experiments were run using an Intel Xeon E5620 CPU.

Table 3 shows a size comparison of each representation by the only common size metric: operators in the logical circuit that each decision diagram induces. We can see that WPBDDs have 60% less logical operators than the corresponding OBDDs on average at both stages of compilation, reducing inference time and system requirements considerably. Also, a WPBDD is reduced by 15% on average by applying the collapse rule when compiling CPTs, and 6% reduction on average with fully compiled networks. This statistic is fully determined by the amount of local structure in the BN and the ordering used during compilation, and can greatly be improved upon utilizing techniques as dynamic compilation in the future.

Observe that there is a close relation between the size of OBDDs and SDD^r s, as mentioned in [43]. We can see that the size of each SDD^r is marginally smaller than its corresponding OBDD in Table 3. We assume that this is because SDDs have multi-valued logical-OR operators, which allow for more concise representations. SDDs consist of binary logical-AND, and n -ary logical-OR operators. We have included OR operators in size computations as $n - 1$ binary logical-OR operators.

¹Available at <http://vlsi.colorado.edu/~fabio/>

²Available at <http://reasoning.cs.ucla.edu/sdd/>

Bayesian Network	Number of operators per CPT						Total number of operators					
	WPBDD	WPBDD ^{nc}	OBDD	ZBDD	SDD ^r	SDD	WPBDD	WPBDD ^{nc}	OBDD	ZBDD	SDD ^r	SDD
example	11	19	48	54	33	49	9	15	39	30	27	49
cancer	80	80	195	747	135	292	66	66	159	324	144	362
earthquake	80	80	195	747	135	292	66	66	159	324	144	362
survey	147	147	354	1671	270	510	132	132	312	825	291	819
asia	131	136	321	1473	231	425	136	136	321	564	306	765
student_farm	231	252	591	3522	441	747	459	465	1098	2007	1080	2135
sachs	747	759	1788	20928	1611	3176	630	630	1602	13164	1581	4881
poker	731	1128	2373	6774	2289	4219	912	1095	2370	5394	2355	6082
child	1011	1099	2739	26088	2385	4663	2934	2988	7872	21861	7857	16030
carpo	1257	1386	3264	77049	2574	4514	2499	3015	7179	16233	7164	13405
powerplant	1414	1558	3795	86184	3141	6382	4158	4362	11043	36276	11025	26662
alarm	1553	1701	3795	41688	3312	6183	3832	4294	10008	19227	9993	35004
hepar2	8371	8467	18528	1593654	16101	29584	56574	56871	142806	2658189	142791	188453
weeduk	29196	30543	70863	13762830	69135	170330	32114	34832	109734	23840949	109455	-
fungiuk	68430	81435	295458	88536714	287940	250500	234715	251739	733551	95350635	727209	-
win95pts	1948	2020	4722	122244	3822	6571	312191	320183	743631	1415904	743619	426550
insurance	2776	3171	7455	106959	6810	11836	468514	474682	1263420	5465610	1263393	1731502
pathfinder	23366	40838	240915	6158832	237549	295272	1899921	2135180	5732988	29435283	5732976	2287777
hailfinder	8909	9520	25371	457860	23865	35707	11141550	11270157	31493220	137548032	31493157	10508499
water	23498	25458	56268	4776201	53979	81614	18460995	18561108	44005977	*	-	-
mildew	139386	833215	1482432	56435169	-	832829	21698281	22322743	71621388	*	-	-
andes	4592	5546	11667	1125528	9192	13809	-	-	-	-	-	-
mainuk	232466	240710	577866	216698319	567135	*	-	-	*	*	-	-
barley	244447	252935	649902	581976891	637101	*	-	-	*	-	-	-
munin	145258	179744	435102	363002097	415134	723740	-	-	-	-	-	-
munin4	145373	180658	432672	345413400	413277	733401	-	-	*	-	-	-
munin3	140318	172843	420537	327439035	400932	702203	*	*	-	-	-	-
diabetes	186630	288075	707679	111194403	696837	984197	-	-	-	-	-	-
pigs	19953	22482	49872	6976569	44127	68288	-	-	*	-	-	-
munin1	25415	32188	74613	7756470	71538	131514	-	-	-	-	-	*
link	22869	27249	60279	12387426	54291	85356	-	-	-	-	-	-
munin2	133493	165853	398550	323654088	379137	670702	-	-	-	-	-	-

Table 3 Number of arithmetic operators in intermediate and resulting decision diagrams
(symbols - and * indicate compilation failure due to memory or an one hour time limitation, respectively)

Bayesian Network	CPT compile time						Conjoin compile time					
	WPBDD	WPBDD ^{nc}	OBDD	ZBDD	SDD ^r	SDD	WPBDD	WPBDD ^{nc}	OBDD	ZBDD	SDD ^r	SDD
example	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
cancer	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
earthquake	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
survey	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
asia	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
student_farm	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.001	0.000	0.000	0.000	0.001
sachs	0.000	0.000	0.000	0.000	0.002	0.007	0.001	0.001	0.000	0.001	0.000	0.012
poker	0.001	0.001	0.002	0.002	0.015	0.030	0.002	0.004	0.000	0.000	0.001	0.016
child	0.000	0.000	0.001	0.000	0.004	0.009	0.004	0.004	0.002	0.007	0.010	0.041
carpo	0.001	0.001	0.003	0.002	0.006	0.008	0.005	0.006	0.004	0.012	0.018	0.033
powerplant	0.001	0.001	0.002	0.002	0.005	0.006	0.007	0.008	0.003	0.025	0.018	0.045
alarm	0.001	0.001	0.002	0.002	0.008	0.018	0.006	0.007	0.003	0.008	0.015	0.059
hepar2	0.006	0.005	0.016	0.063	0.055	0.097	0.072	0.072	0.201	7.045	1.192	3.047
weeduk	0.815	0.817	0.249	1.178	1.615	2512.160	0.084	0.112	0.013	3.988	0.077	(92.86%)
fungiuk	1.799	1.803	1.027	11.011	7.276	179.470	126.701	356.738	0.264	1285.111	1.455	(14.29%)
win95pts	0.003	0.003	0.003	0.005	0.014	0.016	0.129	0.132	1.257	4.203	6.293	0.882
insurance	0.002	0.002	0.004	0.005	0.020	0.022	0.201	0.211	0.366	3.070	1.946	2.485
pathfinder	0.685	0.678	7.840	8.066	30.405	2.226	1.333	1.501	13.974	101.344	67.038	22.888
hailfinder	0.011	0.011	0.018	0.035	0.105	0.177	5.556	5.546	17.409	567.997	71.913	9.464
water	0.107	0.107	0.090	0.253	0.764	4.922	44.400	45.883	55.055	(32.26%)	(96.77%)	(32.26%)
mildew	275.356	275.445	129.703	134.639	(11.43%)	1146.805	304.752	2062.466	160.524	(5.88%)	(0.00%)	(2.94%)
andes	0.006	0.006	0.007	0.032	0.027	0.041	(94.52%)	(94.52%)	(21.92%)	(21.00%)	(20.55%)	(21.92%)
mainuk	7.503	7.528	2.272	27.709	20.179	(4.17%)	(95.744%)	(95.74%)	(44.68%)	(8.51%)	(42.55%)	(0.00%)
barley	7.682	7.683	2.247	77.280	27.284	(4.17%)	(95.74%)	(93.62%)	(44.68%)	(2.13%)	(42.55%)	(0.00%)
munin	0.193	0.190	0.458	30.999	21.165	2.568	(75.67%)	(75.67%)	(1.83%)	(1.35%)	(1.73%)	(2.40%)
munin4	0.186	0.189	0.438	28.238	21.256	2.560	(78.98%)	(78.98%)	(1.54%)	(1.06%)	(1.35%)	(1.45%)
munin3	0.165	0.173	0.395	26.503	23.430	2.225	(74.23%)	(74.23%)	(2.50%)	(1.25%)	(1.63%)	(2.40%)
diabetes	3.142	3.137	6.479	13.272	106.197	11.352	(54.85%)	(54.85%)	(1.21%)	(0.73%)	(0.97%)	(1.21%)
pigs	0.016	0.016	0.036	0.228	0.248	0.114	(94.09%)	(94.09%)	(7.05%)	(6.59%)	(6.59%)	(8.18%)
munin1	0.037	0.036	0.101	0.322	0.998	0.416	(89.73%)	(89.73%)	(14.05%)	(11.89%)	(11.89%)	(12.97%)
link	0.040	0.040	0.070	0.621	0.512	0.403	(89.200%)	(89.20%)	(3.51%)	(2.95%)	(3.09%)	(3.23%)
munin2	0.162	0.162	0.400	28.014	19.650	2.255	(78.94%)	(78.94%)	(1.30%)	(0.60%)	(1.20%)	(1.098%)

Table 4 Compilation time in seconds.

(In case of compilation failure, the percentage of successfully conjoined/processed variables is shown, of which the reason is documented in Table 3)

Bayesian Network	Queries	WPBDD		OBDD			Dlib	
		T	T_{total}	T	T_{total}	S	T	S
example	17	0.000	0.000	0.000	0.000	1.20	0.001	34.70
cancer	714	0.001	0.001	0.002	0.002	1.57	0.138	96.33
earthquake	714	0.001	0.001	0.002	0.002	1.50	0.139	96.53
survey	4448	0.014	0.014	0.023	0.023	1.70	2.244	165.40
asia	18360	0.053	0.053	0.108	0.108	2.03	9.798	185.70
sachs	258324	2.142	2.143	5.323	5.323	2.50	1471.963	687.23
student_farm	1109885	7.346	7.347	17.942	17.943	2.47	2627.214	357.73
poker	145999	2.204	2.209	5.229	5.231	2.37	3574.706	1621.99
child	121647	4.058	4.062	14.603	14.611	3.60	3545.921	873.81
carpo	213257	7.697	7.704	21.340	21.355	2.77	3515.298	456.73
powerplant	457407	21.575	21.584	64.450	64.477	2.99	3393.056	157.27
alarm	94034	5.062	5.07	13.324	13.334	2.63	3547.504	700.77
hepar2	26530	16.419	16.496	50.950	58.011	3.10	3463.509	210.95
weeduk	1051	0.420	1.349	1.608	5.845	3.83	3586.316	8538.88
fungiuk	718	1.911	360.452	6.556	1292.7	3.43	3567.136	1866.32
win95pts	13595	50.586	50.721	135.639	139.845	2.68	3220.060	63.65
insurance	280	1.391	1.604	4.477	7.551	3.22	3571.825	2567.81
pathfinder	765	23.501	25.68	78.277	187.461	3.33	3391.518	144.32
hailfinder	875	133.381	138.938	432.411	1000.43	3.24	2448.653	18.36
water	2221	474.047	520.037	1319.978	1352.33	2.78	*	*
mildew	1312	290.731	2628.64	1401.271	1536.85	3.59	*	*
Avg Speedup						2.69		991.82

Table 5 Inference time in seconds.

In order to evaluate the WMC approach to exact inference with other methods, we have chosen to compare to the Junction tree algorithm using the publicly available Dlib³ C++ library (version 18.18), and the HUGIN⁴ library (through the C++ API version 8.4). We exhaustively go through all possible probabilistic queries. Table 5 and 6 show how much time T spent by each method on an identical set of queries. We have excluded time spent on reading or processing the Bayesian network, as well as creating the join tree, purely focusing on inference time. We went through all possible queries up to `poker`, and limited others by a reasonable amount of time. Time T_{total} indicates the total time spent on compilation and inference, and shows that it occasionally depends on how many queries you intend to answer which language must be chosen to get results faster. In theory, the speedup S of WPBDDs vs other logical representations coincides with the sizes difference of the arithmetic circuits they induce (see Table 3), as inference has linear complexity in the size of induced circuits. This is confirmed with an average speedup of over 2.6x compared to OBDDs. We achieved an average speedup of over 5x compared to HUGIN (Note that we were not able to process all BNs as we used the LITE (free) version, which comes with limitations). We also achieved a staggering speedup compared to the Junction tree algorithm by Dlib, confirmed by an exceptional amount of cache misses reported by cachegrind (Valgrind tool), and other profile information by GNU Gprof and GNU Perf on resource usage. Collectively, compile

³Available at <http://dlib.net/>

⁴<http://www.hugin.com>

Bayesian Network	Queries	WPBDD T	HUGIN	
			T	S
example	17	0.000	0.000	3.865
cancer	714	0.002	0.011	6.209
earthquake	714	0.002	0.012	6.390
survey	4448	0.016	0.085	5.279
asia	14742	0.051	0.407	8.057
sachs	148245	1.347	4.689	3.482
student_farm	608263	4.721	29.264	6.199
poker	185881	3.260	4.929	1.512
Avg Speedup				5.124

Table 6 Inference time in seconds.

and inference results show that WPBDDs make a valuable addition in the field of exact probabilistic inference.

8. Conclusion

To reduce the cost of Bayesian inference through Weighted Model Counting (WMC), we proposed a new canonical language called *Weighted Positive Binary Decision Diagrams* that represent probability distributions more concisely. We have provided theoretical results in addition to practical results on compilation size with regard to 30+ Bayesian networks, where we have seen WPBDD induced logical circuits reduced by 60% on average in comparison to OBDD induced circuits. The introduced reduction rule is responsibly for a 15% reduction on average among compiled CPTs. These results can be improved upon even further in the future by finding a better variable ordering, which is made easier by WPBDDs, as they do not consider probabilities as auxiliary literals, reducing the search space considerably. We have evaluated the cost of inference compared to OBDD induced circuits, yielding a 2.5x speedup on average, and to the Junction tree algorithm, approaching a speedup of 1000x. The language thus gives computational benefits during model counting as well as compilation.

References

References

- [1] J. Pearl, Probabilistic reasoning in intelligent systems: networks of plausible inference, 1988.
- [2] D. Heckerman, J. Breese, Causal independence for probabilistic assessment and inference using Bayesian networks, IEEE Transactions on Systems, Man and Cybernetics 26 (6) (1996) 826–831.
- [3] C. Boutilier, N. Friedman, M. Goldszmidt, D. Koller, Context-specific independence in Bayesian networks, in: Proceedings of the Twelfth international conference on Uncertainty in artificial intelligence, 1996, pp. 115–123.
- [4] F. Jensen, S. Anderson, Approximations in Bayesian belief universe for knowledge based systems, arXiv preprint arXiv:1304.1101.

- [5] F. Bacchus, S. Dalmao, T. Pitassi, DPLL with caching: A new algorithm for #SAT and Bayesian inference, in: *Electronic Colloquium on Computational Complexity*, Vol. 10, 2003.
- [6] M. Chavira, A. Darwiche, On probabilistic inference by Weighted Model Counting, *Artificial Intelligence* 172 (2008) 772–799.
- [7] M. Wachter, R. Haenni, Logical compilation of Bayesian networks with discrete variables, in: *Symbolic and Quantitative Approaches to Reasoning with Uncertainty*, 2007, pp. 536–547.
- [8] J. R. Burch, E. M. Clarke, K. L. McMillan, D. L. Dill, L.-J. Hwang, Symbolic model checking: 10^{20} states and beyond, in: *Proceedings of the Fifth Symposium on Logic in Computer Science*, 1990, pp. 428–439.
- [9] R. D. Shachter, B. D’Ambrosio, B. Del Favero, Symbolic probabilistic inference in belief networks., in: *AAAI*, Vol. 90, 1990, pp. 126–131.
- [10] C. W. Barrett, R. Sebastiani, S. A. Seshia, C. Tinelli, Satisfiability modulo theories, *Handbook of satisfiability* 185 (2009) 825–885.
- [11] F. Somenzi, Cudd: Cu decision diagram package release 3.0.0, Tech. rep., Department of Electrical, Computer, and Energy Engineering, University of Colorado at Boulder (2015).
- [12] A. Choi, D. Kisa, A. Darwiche, Compiling probabilistic graphical models using sentential decision diagrams, in: *Symbolic and Quantitative Approaches to Reasoning with Uncertainty*, 2013, pp. 121–132.
- [13] N. J. Nilsson, Probabilistic logic, *Artificial intelligence* 28 (1986) 71–87.
- [14] J. Y. Halpern, An analysis of first-order logics of probability, *Artificial intelligence* 46 (1990) 311–350.
- [15] D. Poole, First-order probabilistic inference, in: *IJCAI*, Vol. 3, Citeseer, 2003, pp. 985–991.
- [16] F. Bacchus, S. Dalmao, T. Pitassi, Algorithms and complexity results for #SAT and Bayesian inference, in: *In Proceedings of the 44th Symposium on Foundations of Computer Science*, 2003, pp. 340–351.
- [17] T. Walsh, SAT vs CSP, in: *Principles and Practice of Constraint Programming*, 2000, pp. 441–456.
- [18] O. Bailleux, Y. Boufkhad, Efficient CNF encoding of boolean cardinality constraints, in: *Principles and Practice of Constraint Programming*, 2003, pp. 108–122.
- [19] T. Tanjo, N. Tamura, M. Banbara, A compact and efficient SAT-encoding of finite domain CSP, in: *Theory and Applications of Satisfiability Testing-SAT*, 2011, pp. 375–376.
- [20] M. Gavanelli, The log-support encoding of CSP into SAT, in: *Principles and Practice of Constraint Programming*, 2007, pp. 815–822.

- [21] S.-i. Minato, K. Satoh, T. Sato, Compiling Bayesian networks by symbolic probability calculation based on zero-suppressed BDDs., in: IJCAI, 2007, pp. 2550–2555.
- [22] H. Poon, P. Domingos, Sum-product networks: A new deep architecture, in: International Conference on Computer Vision Workshops, 2011, pp. 689–690.
- [23] T. Sang, P. Beame, H. A. Kautz, Performing Bayesian inference by Weighted Model Counting, in: AAAI, Vol. 5, 2005, pp. 475–481.
- [24] M. Bozga, O. Maler, On the representation of probabilities over structured domains, in: Computer Aided Verification, 1999, pp. 261–273.
- [25] N. L. Zhang, D. Poole, On the role of context-specific independence in probabilistic inference, in: Proceedings of the International Joint Conference on Artificial Intelligence, Vol. 85, 1999.
- [26] A. Cano, S. Moral, A. Salmeron, Penniless propagation in join trees, International Journal of Intelligent Systems 15 (2000) 1027–1059.
- [27] A. V. Kozlov, D. Koller, Nonuniform dynamic discretization in hybrid networks, in: Proceedings of the Thirteenth conference on Uncertainty in artificial intelligence, 1997, pp. 314–325.
- [28] A. Cano, M. Gómez-Olmedo, S. Moral, C. B. Pérez-Ariza, Recursive probability trees for Bayesian networks, in: Current Topics in Artificial Intelligence, 2009, pp. 242–251.
- [29] M. Jaeger, Probabilistic decision graphs combining verification and ai techniques for probabilistic inference, International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems 12 (2004) 19–42.
- [30] C. P. Gomes, A. Sabharwal, B. Selman, Model Counting.
- [31] P. Beame, R. Impagliazzo, T. Pitassi, N. Segerlind, Memoization and DPLL: Formula caching proof systems, in: Proceedings of the 18th Conference on Computational Complexity, 2003, pp. 248–259.
- [32] S. M. Majercik, M. L. Littman, Using caching to solve larger probabilistic planning problems, in: AAAI/IAAI, 1998, pp. 954–959.
- [33] E. Fischer, J. A. Makowsky, E. V. Ravve, Counting truth assignments of formulas of bounded tree-width or clique-width, Discrete Applied Mathematics 156 (2008) 511–529.
- [34] P. Sen, A. Deshpande, L. Getoor, Prdb: managing and exploiting rich correlations in probabilistic databases, The VLDB Journal The International Journal on Very Large Data Bases 18 (2009) 1065–1090.
- [35] K. Kersting, B. Ahmadi, S. Natarajan, Counting belief propagation, in: Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence, 2009, pp. 277–284.
- [36] W. Li, P. Poupart, P. Van Beek, Exploiting causal independence using Weighted Model Counting, in: AAAI, 2008, pp. 337–343.

- [37] P. Sen, A. Deshpande, L. Getoor, Read-once functions and query evaluation in probabilistic databases, *Proceedings of the VLDB Endowment* 3 (2010) 1068–1079.
- [38] R. Mateescu, R. Dechter, R. Marinescu, AND/OR Multi-Valued Decision Diagrams (AOMDDs) for graphical models., *J. Artif. Intell. Res.(JAIR)* 33 (2008) 465–519.
- [39] T. D. Nielsen, P.-H. Willemin, F. V. Jensen, U. Kjærulff, Using robdds for inference in bayesian networks with troubleshooting as an example, in: *Proceedings of the Sixteenth conference on Uncertainty in artificial intelligence*, 2000, pp. 426–435.
- [40] E. M. Clarke, K. L. McMillan, X. Zhao, M. Fujita, J. Yang, Spectral transforms for large boolean functions with applications to technology mapping, in: *Proceedings of the 30th Conference on Design Automation*, 1993, pp. 54–60.
- [41] Y.-T. Lai, S. Sastry, Edge-Valued Binary Decision Diagrams for multi-level hierarchical verification, in: *Proceedings of the 29th Design Automation Conference*, 1992, pp. 608–613.
- [42] S. Sanner, D. McAllester, Affine algebraic decision diagrams (AADDs) and their application to structured probabilistic inference, in: *IJCAI*, Vol. 2005, 2005, pp. 1384–1390.
- [43] A. Darwiche, Sdd: A new canonical representation of propositional knowledge bases, in: *Proceedings of the International Joint Conference on Artificial Intelligence*, Vol. 22, 2011, p. 819.
- [44] S.-i. Minato, Zero-suppressed BDDs for set manipulation in combinatorial problems, in: *30th Conference on Design Automation*, 1993, pp. 272–277.
- [45] H. H. Hoos, SAT-encodings, search space structure, and local search performance, in: *IJCAI*, Vol. 99, 1999, pp. 296–303.
- [46] F. M. Brown, *Boolean Reasoning: The Logic of Boolean Equations*, 1990.
- [47] R. E. Bryant, Symbolic Boolean manipulation with ordered binary decision diagrams, *ACM Computing Surveys* 24 (1992) 293–318.
- [48] R. E. Bryant, Graph-based algorithms for Boolean function manipulation, *Transactions on Computers* 100 (1986) 677–691.
- [49] K. S. Brace, R. L. Rudell, R. E. Bryant, Efficient implementation of a BDD package, in: *Proceedings of the 27th design automation conference*, 1991, pp. 40–45.

Appendices

Theorem 1. [46] Shannon's expansion allows a Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ defined over variables X , to be written in terms of its inputs:

$$f = x \wedge f_{|x} \vee \bar{x} \wedge f_{|\bar{x}},$$

with $x \in X$ and where $f_{|x}$ is called the positive cofactor of f with respect to x , and $f_{|\bar{x}}$ the negative cofactor. Applying Shannon's theorem is known as an (additive) decomposition step. The decomposition of f is defined as the recursive application of Shannon's expansion theorem to cofactors, removing one variable at a time, until no variables are left.

Proof. See [46]. □

Lemma 1. A positive Shannon expansion allows an encoded Boolean function $\mathcal{E}(f) = f^e$, where f is defined over X , to be written in terms of its inputs:

$$f^e = f^c \wedge \left(x_i \wedge f_{|x_i}^e \vee f_{|\bar{x}_i}^e \right),$$

where $f^e = f^c \wedge f^m$, and $x_i \in \mathcal{A}(x)$, with $x \in X$. The positive cofactor $f_{|x_i}^e$ incorporates implicit conditioning (Definition 2). The negative cofactor $f_{|\bar{x}_i}^e$ and the decomposition of f^e are as defined by Theorem 1.

Proof. We show that the theorem holds by reducing the Shannon expansion to the positive Shannon expansion through equivalence. Function f is defined over variables $X = \{x\}$, and encoded as Boolean function $f^e = f^c \wedge f^m$ given \mathcal{E} , where:

$$\mathcal{A}(x) = \{x_1, \dots, x_n\}, \quad f^c = \left(\bigvee_{x_j \in \mathcal{A}(x)} \bar{x}_j \right) \wedge \left(\bigwedge_{x_k \in \mathcal{A}(x)} \bigwedge_{x_l \in \mathcal{A}(x) \setminus x_k} (\bar{x}_k \vee \bar{x}_l) \right), \quad f^m = 1.$$

It follows from Equation 2 and 3 that f^c has a cardinality of at least two when encoding a non-trivial function, and that f^m essentially depends on a (non-strict) subset of variables that f^c essentially depends on. We therefore chose f^m to be simplistic to make the following reductions more intuitive. Note that the *at-most-once* clauses generated for x_i are subsumed by f^c at the final step.

$$\begin{aligned} f^e &= f^c \wedge f^m \\ &= \underbrace{\left(\bigvee_{x_j \in \mathcal{A}(x)} x_j \right) \wedge \left(\bigwedge_{x_k \in \mathcal{A}(x)} \bigwedge_{x_l \in \mathcal{A}(x) \setminus x_k} (\bar{x}_k \vee \bar{x}_l) \right)}_{f^c} \wedge \underbrace{1}_{f^m} \end{aligned}$$

$$\begin{aligned}
&= x_i \wedge \underbrace{\left(\bigwedge_{x_j \in \mathcal{A}(x) \setminus x_i} \overline{x_j} \right)}_{f_{|x_i}^e} \vee \overline{x_i} \wedge \underbrace{\left(\bigvee_{x_j \in \mathcal{A}(x) \setminus x_i} \overline{x_j} \right) \wedge \left(\bigwedge_{x_k \in \mathcal{A}(x) \setminus x_i} \bigwedge_{x_l \in \mathcal{A}(x) \setminus x_{\{ik\}}} (\overline{x_k} \vee \overline{x_l}) \right)}_{f_{|\overline{x_i}}^e} \\
&\quad \text{Shannon expansion of } f^e \text{ on } x_i. \\
&= x_i \wedge \underbrace{\left(\bigwedge_{x_j \in \mathcal{A}(x) \setminus x_i} \overline{x_j} \right)}_{f_{\parallel x_i}^e} \wedge \underbrace{1 \vee \overline{x_i}}_{f_{\parallel x_i}^e} \wedge \underbrace{\left(\bigvee_{x_j \in \mathcal{A}(x) \setminus x_i} \overline{x_j} \right) \wedge \left(\bigwedge_{x_k \in \mathcal{A}(x) \setminus x_i} \bigwedge_{x_l \in \mathcal{A}(x) \setminus x_{\{ik\}}} (\overline{x_k} \vee \overline{x_l}) \right)}_{f_{|\overline{x_i}}^e} \\
&\quad \text{Positive Shannon expansion of } f^e \text{ on } x_i. \text{ Equivalent by Definition 2, and Equation 7.} \\
&= x_i \wedge \left(\bigwedge_{x_j \in \mathcal{A}(x) \setminus x_i} \overline{x_j} \right) \wedge f_{\parallel x_i}^e \vee \overline{x_i} \wedge f_{|\overline{x_i}}^e \\
&= \left(\bigwedge_{x_j \in \mathcal{A}(x) \setminus x_i} (\overline{x_i} \vee \overline{x_j}) \right) \wedge \left(x_i \wedge f_{\parallel x_i}^e \vee f_{|\overline{x_i}}^e \right) \\
&= f^c \wedge \left(x_i \wedge f_{\parallel x_i}^e \vee f_{|\overline{x_i}}^e \right)
\end{aligned}$$

□

Theorem 2. A reduced positive Shannon expansion allows an encoded Boolean function $\mathcal{E}(f) = f^e$, where f is defined over X , to be written in terms of its inputs under constraints f^c :

$$f^e \models x_i \wedge f_{\parallel x_i}^e \vee f_{|\overline{x_i}}^e,$$

where $f^e = f^c \wedge f^m$, and $x_i \in \mathcal{A}(x)$, with $x \in X$. Cofactors and decomposition are as defined by Lemma 1.

Proof. We first show what models are introduced by one reduced positive Shannon expansion, providing clear implications what models are introduced by n expansions (i.e. a decomposition). We then show that f^c can be factored out of the positive Shannon decomposed function, and used to remove the introduced models, as they are subsumed by f^c .

Let function f and its encoding be defined as provided in the proof of Lemma 1. We will first show that applying one reduced positive Shannon expansion removes at-most-once (AMO) clauses related to the variable we expand.

$$\begin{aligned}
f^e &= f^c \wedge f^m \\
&= \underbrace{\left(\bigvee_{x_j \in \mathcal{A}(x)} x_j \right) \wedge \left(\bigwedge_{x_k \in \mathcal{A}(x)} \bigwedge_{x_l \in \mathcal{A}(x) \setminus x_k} (\overline{x_k} \vee \overline{x_l}) \right)}_{f^c} \wedge \underbrace{1}_{f^m} \\
&= x_i \wedge \underbrace{1 \vee \left(\bigvee_{x_j \in \mathcal{A}(x) \setminus x_i} \overline{x_j} \right)}_{f_{\parallel x_i}^e} \wedge \underbrace{\left(\bigwedge_{x_k \in \mathcal{A}(x) \setminus x_i} \bigwedge_{x_l \in \mathcal{A}(x) \setminus x_{\{ik\}}} (\overline{x_k} \vee \overline{x_l}) \right)}_{f_{|\overline{x_i}}^e} \\
&= f^w
\end{aligned}$$

We have $M_j \models f^e$, with $j = \{1, \dots, n\}$, where $M_j = \{\overline{x_1}, \dots, \overline{x_{j-1}}, x_j, \overline{x_{j+1}}, \dots, \overline{x_n}\}$ and $M_k \models f^w$, with $k = \{1, \dots, (n-1) + 2^{(n-1)}\}$, where $M_k = \{\overline{x_1}, \dots, \overline{x_{k-1}}, x_k, \overline{x_{k+1}}, \dots, \overline{x_n}\}$ for $1 \leq k \leq n$ and $k \neq i$, and $M_k = \{\dots, x_i, \dots\}$ for the remainder. We conclude that $f^e \neq f^w$, because f^e has n models, while f^w has $(n-1) + 2^{(n-1)}$. Expanding x_i has caused any model containing x_i to be true, i.e., model $\{\overline{x_1}, \dots, \overline{x_{i-1}}, x_i, \overline{x_{i+1}}, \dots, \overline{x_n}\}$ has changed to the models $\{\dots, x_i, \dots\}$. It is precisely those additional models in M_k that are not in M_j , which can be removed by AMO clauses created for x_i , i.e.:

$$f^e \equiv \left(\bigwedge_{x_j \in \mathcal{A}(x) \setminus x_i} (\overline{x_i} \vee \overline{x_j}) \right) \wedge f^w.$$

Where one expansion on x_i removes AMO clauses related to x_i , a decomposition clearly removes AMO clauses related to all variables $\mathcal{A}(x)$. This holds regardless of ordering, as a positive Shannon decomposition of f^e is guaranteed to produce an isomorphic representation due to the symmetric nature of the constraints. We can reduce the positive Shannon expansion to its reduced by form factoring out f^c :

$$\begin{aligned} f^w &= f^c \wedge f^m \\ &= \underbrace{\left(\bigvee_{x_j \in \mathcal{A}(x)} \overline{x_j} \right)}_{f^c} \wedge \underbrace{\left(\bigwedge_{x_k \in \mathcal{A}(x)} \bigwedge_{x_l \in \mathcal{A}(x) \setminus x_k} (\overline{x_k} \vee \overline{x_l}) \right)}_{f^m} \wedge 1 \\ &= f^c \wedge \left(\underbrace{x_1 \wedge 1}_{f_{\parallel x_1}^e} \vee \underbrace{\left(\bigvee_{x_j \in \mathcal{A}(x) \setminus x_1} \overline{x_j} \right) \wedge \left(\bigwedge_{x_k \in \mathcal{A}(x) \setminus x_1} \bigwedge_{x_l \in \mathcal{A}(x) \setminus \{x_1, k\}} (\overline{x_k} \vee \overline{x_l}) \right)}_{f_{\mid x_1}^e} \right) \\ &= f^c \wedge (x_1 \vee f^c \wedge (x_2 \wedge \underbrace{1}_{f_{\parallel x_2}^e} \vee \underbrace{\dots}_{f_{\mid x_2}^e})) \\ &= f^c \wedge (x_1 \vee f^c \wedge (x_2 \vee f^c \wedge (x_3 \wedge \underbrace{1}_{f_{\parallel x_3}^e} \vee \underbrace{\dots}_{f_{\mid x_3}^e}))) \\ &= \dots \\ &= f^c \wedge (x_1 \vee (f^c \wedge (x_2 \vee (f^c \wedge \dots x_{n-1} \vee (f^c \wedge x_n)))))) \\ &= f^c \wedge (x_1 \vee \dots \vee x_n) \\ &= f^c \end{aligned}$$

This confirms when using reduced positive Shannon decompositions, equivalence is only achieved under domain closure constraints, as AMO clauses are subsumed by f^c . \square

Proposition 1. *An OBDD representing Boolean function f and an PBDD representing $\mathcal{E}(f)$ induce isomorphic logical circuits under Boolean identity, given an appropriate ordering.*

Proof. Let $\mathcal{E}(f) = f^e$ be a Boolean function, with f a boolean function defined over variables $X = \{x^1, \dots, x^n\}$. We show that an OBDD representing f is equivalent to a PBDD representing $\mathcal{E}(f)$ by comparing their induced logical circuits, under the premis that the collapse rule does not apply distribution, but deletes literals. Note that this comparison requires f to be a Boolean function, and thus each $x \in X$ is mapped to two atoms $\mathcal{A}(x) = \{x_1, x_2\}$. For every Shannon expansion on f , there is an equivalent positive Shannon expansion on f^e (Figure 10).




Figure 10 Expansions

We have equality, because there exists a trivial mapping between the two decomposition types, namely $x = x_1$ and $\bar{x} = x_2$. Observe the implication that the collapse rule can be applied to f^e whenever the delete rule can be applied to f . Under our pretense we can infer that the OBDD and PBDD induce equivalent circuits, because they are isomorphic due to a trivial mapping for precisely those orderings where, for each $x \in X$, atoms $\mathcal{A}(x) = \{x_1, x_2\}$ are placed adjacent in the order, e.g., $\mathcal{A}(x^1) < \dots < \mathcal{A}(x^n)$ and for each $x \in X$ we have partial orders $x_1 < x_2$. The delete rule alters the circuit if cofactors are equal by applying the distributive law and identity, e.g., $x \wedge f \vee \bar{x} \wedge f = (x \vee \bar{x}) \wedge f = f$. The collapse rule forgoes that last step. We therefore conclude that the OBDD and PBDD induce equivalent circuits under Boolean identity. \square

Proposition 2. *Given an ordering on $\mathcal{A}(X)$, the size of PBDD φ is less than the size of OBDD ψ when both representing $\mathcal{E}(f) = f^e$, where f is defined over variables X .*

Proof. In the case where f is not Boolean, we will show that PBDD φ is always smaller than OBDD ψ , when both representing $\mathcal{E}(f) = f^e$. Consider some $x \in X$, for which we find atoms $\mathcal{A}(x) = \{x_1, x_2, x_3, \dots, x_n\}$ adjacent in the ordering, where n is the domain size of x . OBDD ψ will contain the subfunction shown in Figure 11.




Figure 11 Constraints in OBDD ψ

Here, j, k, l and m are subfunctions that essentially depend on $\{x_4, \dots, x_n\}$. We can transition from the OBDD induced logical circuit above to the corresponding logical circuit induced by a PBDD, by removing the bold edges, literals and operators. Shannon expansions that have indicator 1 at their root can be removed, which coincides with implicit conditioning. More generally, this allows us to remove $\sum_{i=1}^n (i - 1)$ nodes from the OBDD. Additionally, the implicates for negative cofactors are removed that are marked by indicator 2.

Removing the implicate for the negative cofactor, and implicit conditioning, contribute to reducing the size of induced logical circuits by removing operators while maintaining equivalence. The extent to which is lower bounded by:

$$\sum_{x \in X} \underbrace{n}_{\text{negative cofactor}} + \underbrace{\sum_{i=1}^n (i - 1) * 3}_{\text{implicit conditioning}},$$

where we sum over every $x \in X$, with n the domain size of x , i.e., $|\mathcal{A}(x)|$. This increases when atoms $\mathcal{A}(x)$ are not adjacent in the ordering, and is multiplied by the number of distinct subfunctions in OBDD ψ that depend on $\mathcal{A}(x)$. Furthermore, we are guaranteed to encounter each $x \in \mathcal{A}(X)$ in an OBDD along every path from root to the *true* terminal, which serves as an upper bound regarding PBDDs. The collapse rule can effectively reduce the size of the representation if the structure of f^e allows it, by reducing the number of nodes from $|\mathcal{A}(pa(x))|$ to $|pa(x)|$ for every subfunction, provided that literals $\mathcal{A}(y)$, with $y \in pa(x)$, for each parent are adjacent in the ordering. \square